Detrending the long-term stellar activity and the systematics of the *Kepler* data with a non-parametric approach.

C. Danielski

Dept. of Physics & Astronomy, University College London, Gower Street, WC1E 6BT, UK

camilla@star.ucl.ac.uk

T. Kacprzak, G. Tinetti

Dept. of Physics & Astronomy, University College London, Gower Street, WC1E 6BT, UK

Received _____; accepted _____

Abstract

The NASA Kepler mission is delivering groundbreaking results, with an increasing number of Earth-sized and moon-sized objects been discovered. A high photometric precision can be reached only through a thorough removal of the stellar activity and the instrumental systematics. We have explored here the possibility of using non-parametric methods to analyse the Simple Aperture Photometry data observed by the *Kepler* mission. We focused on a sample of stellar light curves with different effective temperatures and flux modulations, and we found that Gaussian Processes-based techniques can very effectively correct the instrumental systematics along with the long-term stellar activity. Our method can disentangle astrophysical features (*events*), such as planetary transits, flares or general sudden variations in the intensity, from the star signal and it is very efficient as it requires only a few training iterations of the Gaussian Process model. The results obtained show the potential of our method to isolate the main *events* in the light curves for both *Kepler* long cadence and short cadence data (i.e. integration time of 29.4 min and 58.9 s respectively). We tested our approach on the star Kepler-19, finding that the transit depth of its planetary companion is consistent at 1-σ with the one published in the literature.

Subject headings: techniques: photometric, methods: data analysis, planets and satellites: atmospheres,

1. Introduction

The NASA *Kepler* mission (Borucki et al 2010, 2011) has revolutionized the exoplanet-field, providing very accurate observations of thousands of stars. Since its launch in March 2009, a total of 122 planets and more than 2700 planetary candidates have been detected. Those planetary candidates span a vast range of sizes, temperatures, transit periods and host-star types. With its unprecedented photometric precision *Kepler* is capable of groundbreaking discoveries like the detection of the moon-sized planet Kepler-37b (Barclay et al. 2013) or the recent finding of the two super-Earth sized planets in the habitable zone around Kepler-62 (Borucki et al. 2013). Such discoveries suggested that small worlds are common, and, most importantly, are observable today if the host star is bright enough and if we reach the correct photometric precision. To detect even the smallest flux variations, the data reduction techniques adopted are critical. In particular, instrument systematics and stochastic errors need meticulous corrections (Jenkins 2010a). Intense work in this regard has been done by the Kepler team: the Presearch Data Conditioning (PDC) module of the Kepler data analysis pipeline (Jenkins et al. 2010b, Gilliland et al. 2010b, Christiansen et al. 2012) is very efficient at extracting transit signals. However it was not optimized to analyse the low-frequency stellar signal (Murphy 2012) and some instrument systematics may still appear in the residuals after the correction. Alternative ways to solve these hurdles have been suggested: for example García et al. (2011) discussed the process of correction of Kepler time-series focusing more on asteroseismic applications. Smith et al. (2012) applied a Bayesian Maximum A Posteriori approach to model the instrumental systematics by correlating multiple non-active targets. All these approaches either apply parametric corrections - which potentially may distort the signal and inject supplementary noise - or use a limited model class (e.g. linear) that may not catch the real underlying trend. The problem can be mitigated by adopting a non-parametric approach.
Non-parametric techniques have been used in a variety of astronomical contexts (e.g. Seikel

et al. 2012, Chapman et al. 2013). In the exoplanet field in particular, these methods have been adopted by e.g. Thatte et al. 2010, Gibson et al. 2012a, Ford et al. 2012, Waldmann 2012, 2013, Waldmann et al. 2013a. Among these the Gaussian Process (GP) method for regression (Rasmussen & Williams, 2006), is used widely in the machine learning field.

Here we adopted a technique which uses GP to analyse the *Kepler* raw stellar light curves to extract the temporal *events* intrinsic to the stellar system, such as planetary transit or flares. The study of the intrinsic stellar variability will be discussed in future work.

2. Method

2.1. *An introduction to Gaussian Processes*

A Gaussian Process (GP) is a stochastic process (a set of random variables), any finite number of which have a joint Gaussian distribution, (Rasmussen & Williams, 2006). A GP is completely specified by its mean function $m(\mathbf{x})$ and its covariance function $K(\mathbf{x_1}, \mathbf{x_2})$ (also often called the *kernel* function). GP can be thought of as a probability distribution over function space, defined by the first and second statistical moments, the mean $m(\mathbf{x})$ and the covariance $K(\mathbf{x}_1, \mathbf{x}_2)$ between any two points $\mathbf{x}_1, \mathbf{x}_2$. Hence the covariance and mean functions can be considered as a prior on the function space, which is parametrized by so-called *hyperparameters*.

Consider a function $\mathbf{f}$ drawn from a GP. We obtain a finite number i of noisy observed data y_i of this function at a set of points $\mathbf{x}_i$, which can be a vector with some N_d dimensions. We have a prior knowledge about the properties of the function - we choose and specify the covariance and mean functions. Thus, the covariance function is a prior on the function

space

$$p(\mathbf{f}|\mathbf{X}) = \mathcal{N}(\mathbf{m}, \mathbf{K}) \tag{1}$$

where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_2, ..., \mathbf{x}_n]^\top$ is a collection of points, $\mathbf{m} = m(\mathbf{X})$ and $\mathbf{K} = K(\mathbf{X}, \mathbf{X})$ are the mean and covariance functions. A set of observations $\mathbf{y}$ contains additive, white Gaussian noise with standard deviation σ :

$$p(\mathbf{y}|\mathbf{f}) = \mathcal{N}\left(\mathbf{f}, \sigma^2\mathbf{I}\right). \tag{2}$$

Integrating over $\mathbf{f}$ gives the marginal distribution on $\mathbf{y}$

$$p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f})d\mathbf{f} = \mathcal{N}[\mathbf{y}|\mathbf{m}, \tilde{\mathbf{K}}], \tag{3}$$

where $\tilde{\mathbf{K}} = \mathbf{K} + \sigma^2\mathbf{I}$, and $\mathbf{I}$ is an identity matrix.

Joint probability of a set of observations $\mathbf{y}$ and test points $\mathbf{f}_*$ is

$$p\left(\begin{bmatrix}\mathbf{y}\\ \mathbf{f}_*\end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix}\mathbf{m}\\ \mathbf{m}_*\end{bmatrix}, \begin{bmatrix}\tilde{\mathbf{K}} & \mathbf{k}\\ \mathbf{k}^\top & \mathbf{c}\end{bmatrix}\right) \tag{4}$$

where $\mathbf{m}_* = m(\mathbf{X}_*)$ and $\mathbf{k} = K(\mathbf{X}, \mathbf{X}_*)$, $\mathbf{c} = K(\mathbf{X}_*, \mathbf{X}_*)$.

Conditioning on observed values of $\mathbf{y}$ allows us to calculate the conditional probability for $\mathbf{f}_*$, which will be a Gaussian distribution:

$$p(\mathbf{f}_*|\mathbf{y}, \mathbf{X}, \mathbf{X}_*) = \mathcal{N}(\mathbf{m}_* + \mathbf{k}^\top\tilde{\mathbf{K}}^{-1}(\mathbf{y} - \mathbf{m}), \mathbf{c} - \mathbf{k}^\top\tilde{\mathbf{K}}^{-1}\mathbf{k}). \tag{5}$$

The prior on the mean function is often chosen as zero when the data is conveniently normalized. The functional form for the covariance is chosen to match the prior beliefs about the statistical properties of the function $\mathbf{f}$. Often there is no theoretical motivation for any particular choice of the covariance function form and its hyperparameters. In that case, the functional form is chosen from a popular set of functions used in the community, for example Radial Basis Function (RBF), polynomial kernel, rational quadratic, etc. The

hyperparameters are then optimised, often using cross - validation of the training data. In this work we use the RBF kernel defined as

$$K(\mathbf{x}_1, \mathbf{x}_2) = \lambda \exp\left(-\frac{1}{2l^2}(\mathbf{x}_1 - \mathbf{x}_2)^\top(\mathbf{x}_1 - \mathbf{x}_2)\right) \tag{6}$$

where λ and l are kernel hyperparameters. We used a zero mean prior.

The variance of conditional probability $\mathbf{f}_*$ depends on how uncertain we are about the value of the function in point $\mathbf{x}_*$. This uncertainty varies with values and proximity of points in the training set, as well as the noise variance σ^2.

2.1.1. Specifications of the statistical model

We describe here the specifications of our statistical model for the observed light curves. We also report the assumptions we made to model the light curves and the *events* (§2.2).

1. The true luminosity of a star can be estimated through a covariance function, which is of the form expressed in in Eqn (6). The true physical process which governs the luminosity of the star is unknown to us, and it is not currently our concern. We found that the RBF kernel is a good model. To find the hyperparameters of the RBF kernel, we used a optimization procedure implemented in the GPML package, based on data cross-validation (see GPML documentation[1] for details).

2. The observed values of stellar luminosity are contaminated with additive Gaussian noise with some standard deviation, which can be roughly estimated from the data. Although this is an important parameter, a small mis-estimation of it does not change the results significantly.

[1] http://www.gaussianprocess.org/gpml/code/matlab/doc/

3. Except for Gaussian noise, the instrument may introduce a systematic variation in the observed star luminosity, for example a periodic oscillation or more (see §2.2). If this effect is present we do not treat it separately from the star luminosity but used a GP to model both simultaneously.

4. The raw data show *events* which are not governed by the same physical process directly responsible for the stellar variation of the luminosity. We identified as *events* the data points whose luminosity lay beyond the 3σ threshold of the GP probability distribution on the luminosity function (Eqn 5). We specified a parameter t^e_{max} [days], as being the maximum duration of an *event* expected in the data. We removed these data points and a range of neighbouring points, which lay within t^e_{max} range of the identified *event* points. Consequently we re-trained the GP to model the stellar luminosity without the influence of these outliers (see §2.2, Fig. 1).

5. The *events* have a maximum duration t^e_{max} which is smaller than the typical covariance length of the star/systematic luminosity variation. To enforce this assumption, during the GP hyperparameters optimisation step, we did not allow the covariance function parameter l to be smaller than 1.5 t^e_{max}. This heuristic approach worked well with all the data analysed.

   If the maximum effective Fourier frequency of the sum of the star luminosity function and the instrumental systematic is greater than the maximum effective Fourier frequency of this enforced kernel function, this approach can lead to an error in estimation of the parameters of *events* of interest. As this spurious signal has a frequency range which is higher than the *event* signal, other methods can be used to correct for it after subtracting the main stellar trend. However, for all the stars analysed with our GP method, we did not find this to be an issue.

   For example, if this high frequency part of the stellar luminosity is roughly a periodic

feature, another GP with a periodic covariance function can be used to model it and remove it on the residual level. The periodic covariance functions frequently used in Machine Learning community are of the form $\sigma_f^2 \exp(-\frac{2}{l^2}\sin^2[\frac{\omega}{2\pi}(\mathbf{x}_1 - \mathbf{x}_2)])$ where ω controls the period of the oscillatory feature. We leave this investigation to future work.

6. We used a single GP to model each time-series. GP inference requires an inverse of a matrix of size N^2, where N is the number of training points, and is a $\mathcal{O}(N^3)$ process. On a standard desktop computer, this is feasible in reasonable time for data sets with less than 10 000 points. Data sets we analysed had fewer points. However, if more points are needed in the time-series, there exist large scale implementations of GP (see Quionero-Candela 2005 for reviews). Also it is possible to split the GP into overlapping parts, analyse them separately and then join them together, if the length of the parts is much greater than the width of the covariance function used.

2.2. *Data analysis*

The aim of our analysis is to identify clusters of outliers that are due to the astrophysical signatures intrinsic to the stellar systems, such as flares, planetary transits, occultations or general sudden variations in the intensity. Hence we investigated a selection of stars belonging to the Kepler Objects of Interest (KOI) catalog, with different effective temperatures (T_{eff}). For each star in the sample we analysed the respective long cadence (LC) Simple Aperture Photometry (SAP) light curves observed in the quarters "Q", reported in Table 1, with the *Kepler Space Telescope*. All the time-series under study exhibit instrumental perturbations such as data-breaks, discontinuities, flux jumps and drifts (García et al. 2011). The data-breaks are generally due to the monthly downloads of stored science and engineering information, but they are also due to the space telescope

safe mode and loss-of-pointing. Normally the safe mode is responsible for the longest gaps in the data, followed by a sudden exponential increase/decrease in flux as the telescope returns to its nominal temperature and focus. Unexpected jumps in flux can also be due to pointing tweaks and pixel sensitivity drops (Jenkins et al. 2010a). These instrumental effects are in part corrected in the PDC flux, however it has been shown that in some cases PDC-flux data may modify the astrophysical signals (i.e. low-frequency stellar signals such as the ones produced by long/short lived star-spots) not directly related to exoplanetary transits (García et al 2011, Murphy 2012). Note that the SAP-flux data also present an absolute calibration between long cadence (LC) and short-cadence (SC) data. The only difference is that SC pixels, co-added in the Science Data Accumulator, are read once every 58.9 s (9 exposures of the SC versus 270 exposures of the LC data with a total integration time of 29.4 min). SC data collection in fact does not require the hardware to do anything different than what it does for LC data collection and the target for SC and LC is always the same (J. E. Van Cleve & D. A. Caldwell 2009). Here we chose to use the SAP-flux data (henceforth referred to us the *"raw",* data). Each Quarter was analysed separately.

Before proceeding with our analysis we normalized each time-series. Then we removed on average ∼50-100 points, equivalent to ∼1.04-2.08 days, around systematic gaps longer than 0.5 BJD (Barycentric Julian date). Our choice of cutting the data at the extremities of these gaps does not affect the purpose of our analysis and it is a luxury we can afford because of the very large number of data points in the dataset. Moreover, this conservative approach prevents the addition of further systematic trends to our GPs model. The number of data points to cut was chosen in relation to every single light curve.

We applied GPs described in §2.1.1 to the normalized time-series the to predict a model of the flux as a function of time t with mean $m(t)$ and standard deviation σ. Consequently we located the outliers greater than 3σ and, by defining the maximum duration t^e_{max} of the

intensity fluctuations we are looking for in the generative model (see §2.1.1, point 4), we isolated a number of data points adjacent to the outliers. To not affect the amplitude of the sought *events*, we temporally removed the clumps of outliers from the light curves and we trained a new GP model on the remaining data only (hereafter *inliers*). This resulted in a second model prediction with mean $\mu(t)$ and standard deviation σ^*. Figure 1 shows how the amplitude of the signature could be altered if the prediction model is applied on the totality of the data instead of only on the *inliers*. We show in Fig. 2 the fit on the 6th quarter *inliers* data of Kepler KIC-2571238 (most known as Kepler-19, Ballard et al. 2011), with relative residuals.

Finally we resampled the predicted mean $\mu(t)$ on the whole LC light curve, outliers included, and computed the residuals. We obtained a de-trended light curve where the main stellar contribution is removed non-parametrically and the main astrophysical features are located and conserved (transits and flares, Fig. 3, 4 respectively).

Note that, as previously mentioned, the SAP-LC and -SC data differ only by the integration time, nevertheless the SC data exhibit fewer low-frequency artefacts, reduced errors on frequency and the evident increased sampling rate. The latter in particularly is useful for studies of short lived *events* like flares and transits (Murphy 2012). Assuming that the astrophysical signal contained in the LC and SC data is exactly the same, we can apply the model $M[\mu(t), \sigma^*]$ evaluated on the LC data to the SC data, obtaining high resolution star-detrended residuals (Fig. 5). It could have been possible to run the GPs directly on the SC light curves (working on sub-quarter/monthly-download separately), but that is prohibitively intensive computationally (see §2.1.1, point 6). Also, the standard deviation of the noise on the LC is smaller than for the SC time-series, making it is easier to catch the main star-trend on the LC data without suffering the effects of increased Gaussian noise.

3. Results and discussion

We analysed here the long cadence (LC) data of multiple Kepler Quarters for a sample of stars with different stellar activities and effective temperatures T_{eff} (Tab. 1). We have been able to model the stellar luminosity along with the instrumental systematics, disentangling the astrophysical features from the main star modulations. Such astrophysical features, like transits and flares, are recognisable in the light curve as sudden variations in the intensity. Furthermore we used the Gaussian probability distribution, described by the model $M[\mu(t), \sigma^*]$ (see §2.2), to compute high resolution residuals, by resampling the same model M to the short cadence (SC) data. The resolved LC and SC cadence residuals can then be used to confirm multiplanet systems via precise transit timing variation measurements (e.g. Ford et al 2012, Steffen et al 2013), for example. The process of obtaining the final de-trended residuals took less than 2 minutes on a standard desktop computer.

Note that here we used Gaussian Processes to isolate only the temporal *events* that are not governed by the same physical process directly responsible for the star luminosity variation. We leave the study of the intrinsic stellar variability for future work.

3.1. *Technique validation*

To test the GP technique we modelled, as described in §2.2, the long cadence data, quarters Q 3-6, of Kepler-19 (Ballard et al. 2011). We then resampled the same model M_Q $[\mu(t), \sigma^*]$, retrieved for each different quarter Q, on the respective SC quarter. In this way we obtained LC and SC residuals where the planetary transit of Kepler-19b is isolated. Next we folded the residual time-series over the transit feature and we separately fitted each transit light curve with a Mandel & Agol model (Mandel & Agol 2002) using a

Markov chain Monte Carlo (MCMC) technique (Haario et al. 2006). The stellar system and the limb darkening parameters implemented in the fitting function are taken from the discovery paper of Kepler-19b (Ballard et al. 2011). Figure 6 shows the weighted mean depth with respective uncertainties computed for both LC and SC data for each quarter. All the LC and SC measurements are consistent within themselves and consistent with respect to each other. As expected the LC depth uncertainties are larger than the SC ones. Furthermore we calculated the weighted mean of all the SC depth values obtaining $\overline{d} = 5.6161 \cdot 10^{-4} \pm 4.2656 \cdot 10^{-6}$, a value which is consistent with Ballard et al. (2011) at 1-σ.

4. Conclusions

We have studied here new non-parametric techniques to model the *Kepler* Simple Aperture Photometry data. By applying Gaussian Processes-based techniques to a sample of *Kepler* stellar light curves, we disentangled temporal *events*, such as flares, planetary transits and sudden intensity variations, from the long-term stellar modulation. The detrended residuals are sampled both in long cadence and short cadence mode, i.e. with an integration time of 29.4 min and 58.9 s respectively. The last ones are particularly useful to confirm multiplanet systems via precise transit timing variation evaluations.

Our method is innovative as it allows us to model all the instrumental systematics and the stellar signal in a completly non-parametric way, without distorting the signal or injecting noise as may happen with a simple least square fit; moreover it provides an absolute calibration within the *events* that reside in the same quarter.

We tested our method on a known star, Kepler-19, by fitting a transit model light curve to the residuals with a Markov Chain Monte Carlo approach. The measured depth value is consistent with the one published in the discovery paper, validating the functionality of our

technique.

KIC	T_{eff} [K]	R [$R_\odot$]	E(B-V)	Kep_{mag}	Quarter	Flag
3835670	5642	1.790	0.115	13.397	3, 4,5,6	Planetary candidate
2571238	5541	0.850	0.038	11.898	3,4,5,6	Exoplanet (Kepler-19b)
7700622	4789	0.649	0.062	12.968	6,7,8	Planetary candidate
3128793	4668	0.637	0.062	14.633	9,11,12,13	Planetary candidate
7603200	3900	0.612	–	12.925	6,7,8,9,10	Planetary candidate
7907423	3803	0.492	0.032	15.234	11,12,13	Planetary candidate

Table 1: Table of the targets. We report here the effective temperature (T_{eff}), the stellar radius ($R_\odot$), the stellar colour (E(B-V)), the Kepler magnitude (K_{mag}) and flag as in the KOI catalogue. For each target we analysed each quarter Q.

Acknowledgments: We thank Ingo P. Waldmann for his helpful suggestions and useful discussions.

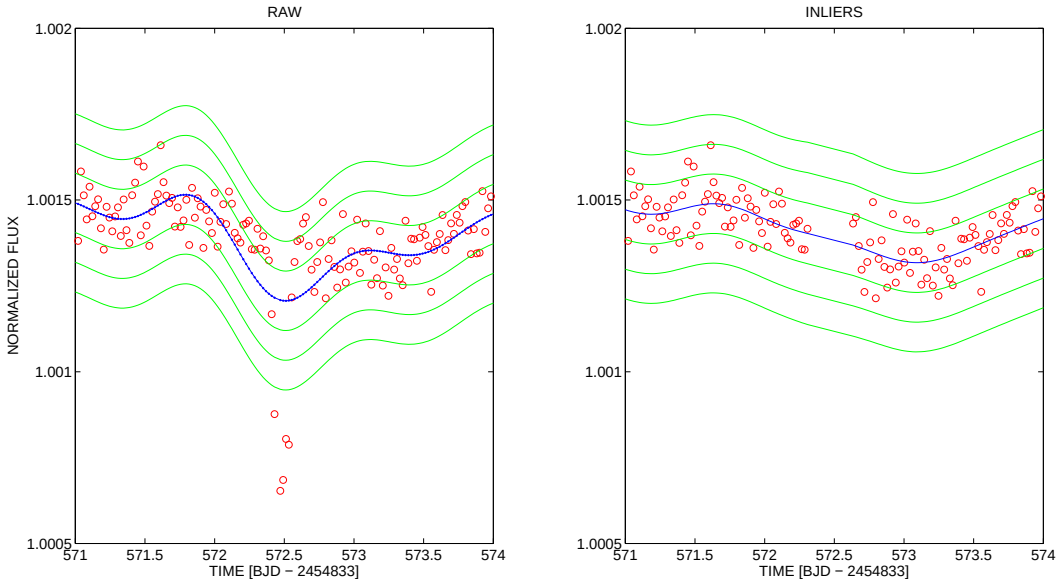

Fig. 1.— GP model prediction on the totality of the data (*raw* - *left*) and after the removal of the clusters of outliers (*inliers* - *right*). The data are plotted in red, while the model and the 1,2,3 σ confidence limits are represented by the blue dotted line and green lines respectively.

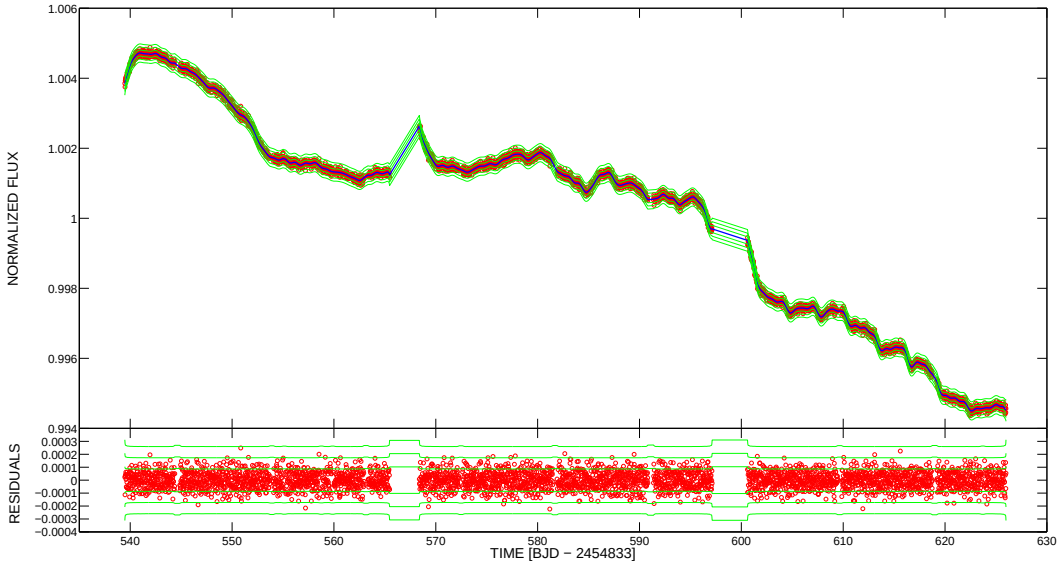

Fig. 2.— GP model prediction (*top*) and residuals (*bottom*) on the long-cadence *inliers* of Kepler-19 (Q-6). The data are plotted in red, while the model and the 1,2,3 σ^* confidence limits are represented by the blue dotted line and green lines respectively.

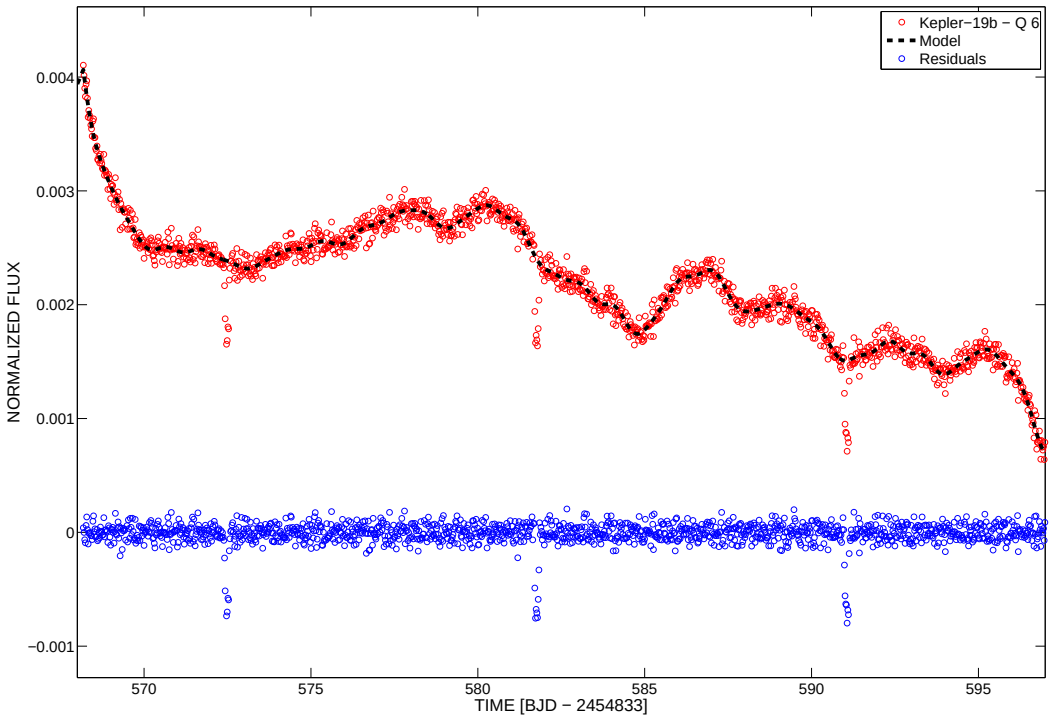

Fig. 3.— GP model (dotted line) applied to the 6th quarter long cadence data (red circles) of Kepler-19 and the respective residuals (blue circles). The raw data are shifted for clarity.

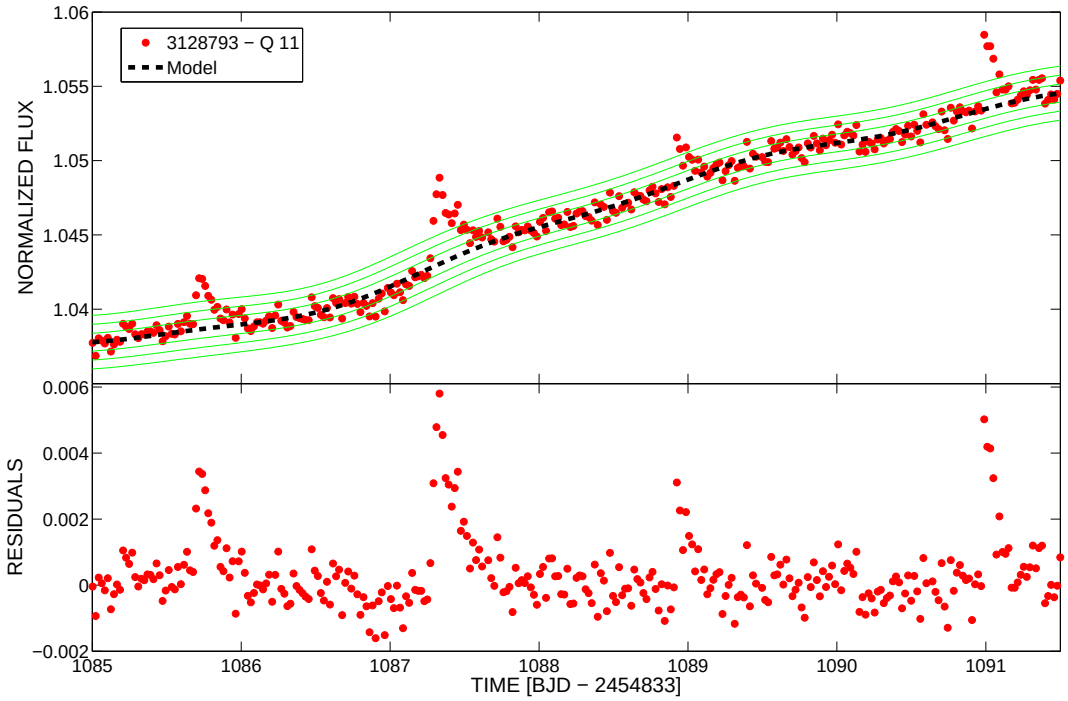

Fig. 4.— Zoom-in on stellar flares of the long cadence observations of KIC 3128793 (Q11). In the top panel we show the fit to the data while in the bottom panel the residuals after the GP detrend are shown. The data are plotted in red, while the model and the 1,2,3 σ^* confidence limits are represented by the blue dotted line and green lines respectively.

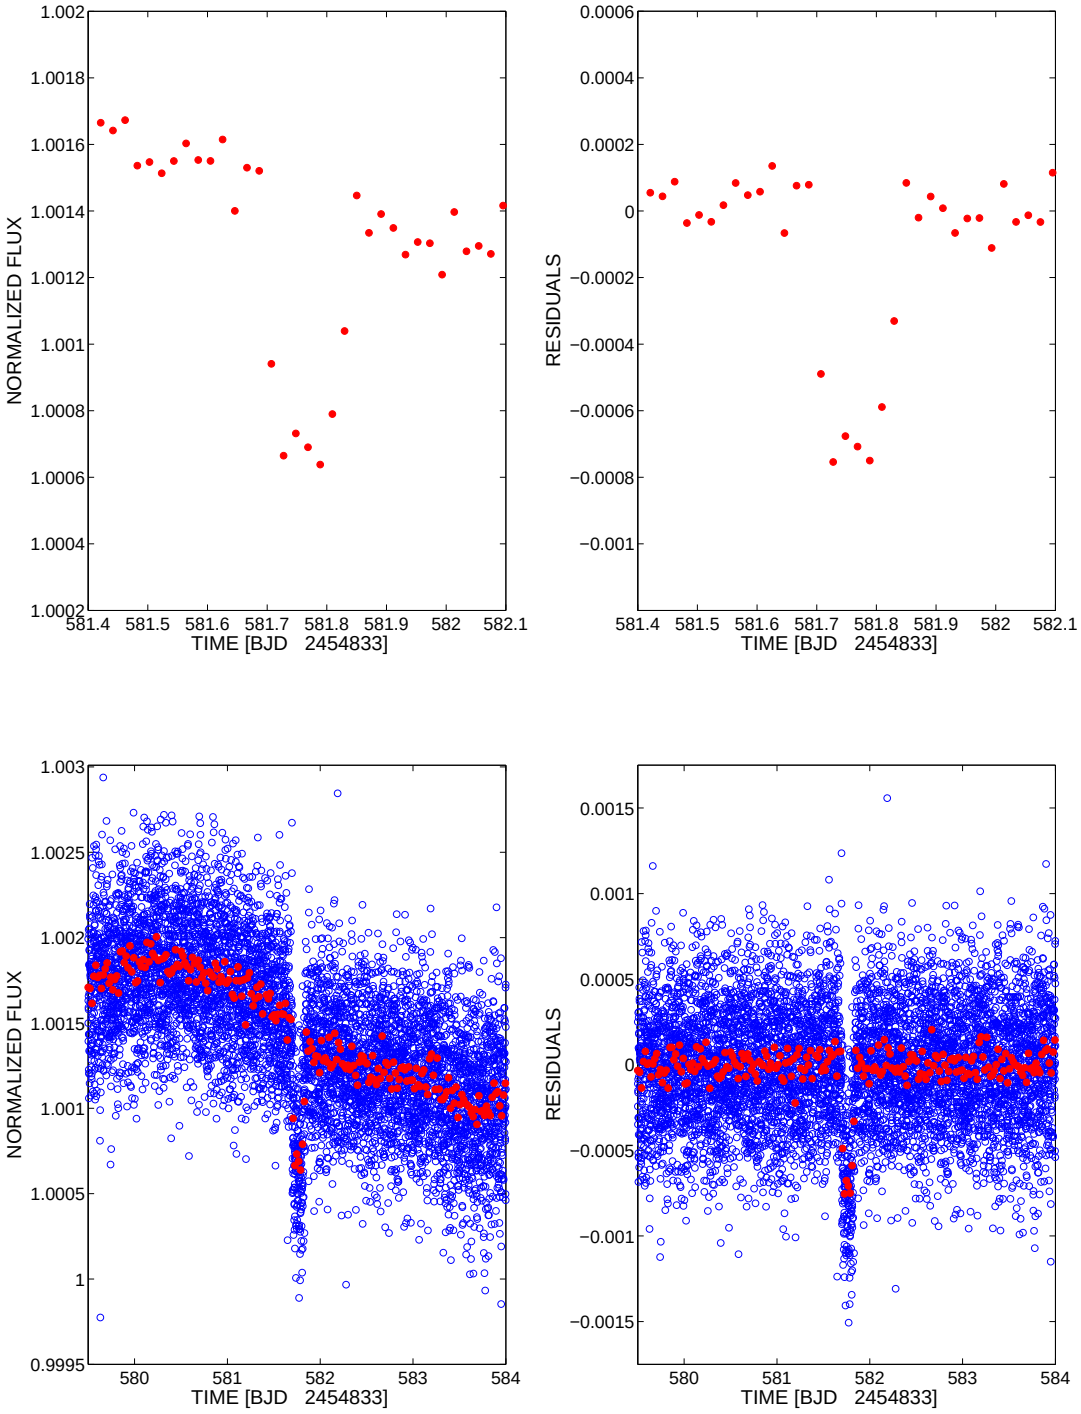

Fig. 5.— Top panel: Zoom-in of a transit of Kepler-19b (Q6) observed at long cadence before (*left*) and after (*right*) the GP detrend technique. Bottom panel: Zoom-in of the same transit of Kepler-19b (Q6) with the application of the GP model, retrieved for the long cadence observations (red dots), to the short cadence data (blue circles).

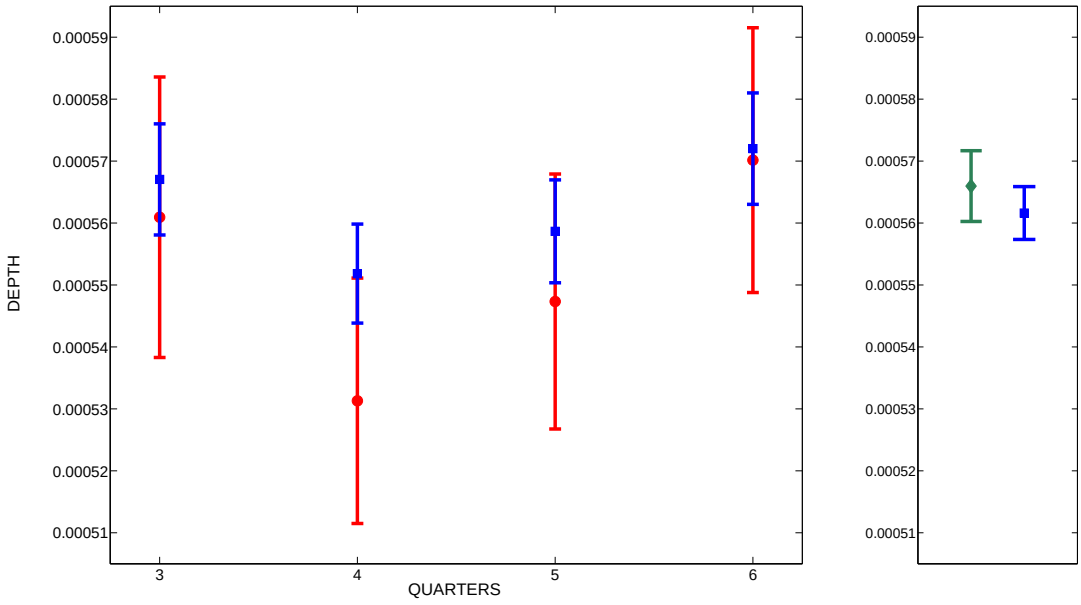

Fig. 6.— *Left:* Mean transit depth values and uncertainties for quarters Q 3,4,5,6 computed for the long cadence data (red circles) and short cadence data (blue squares). *Right:* Mean transit depth of all the quarters (blue square) – short cadence data – compared with the transit depth by Ballard et al. 2011 (green diamond).

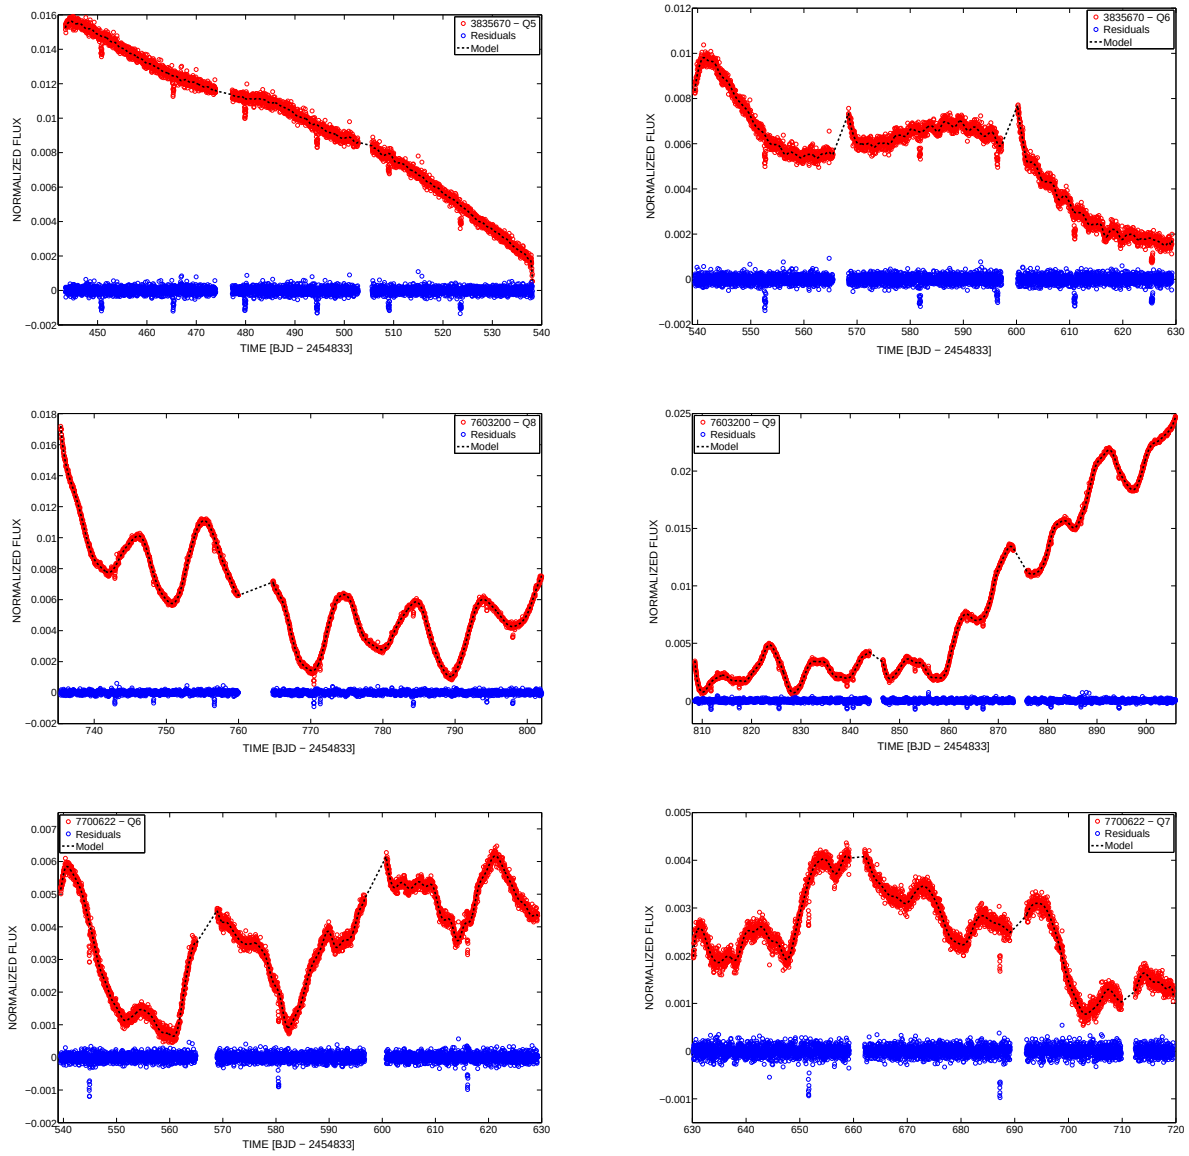

Fig. 7.— Fit to the raw data (red circles) and respective residuals (blue circles) after the GP detrend (dotted line). We plotted a couple of quarters for each analysed planetary candidate (from top to bottom: KIC 3835670 (Q 5,6); KIC 7603200 (Q 8,9); KIC 7700622 (Q 6,7)). The raw data are shifted for clarity.

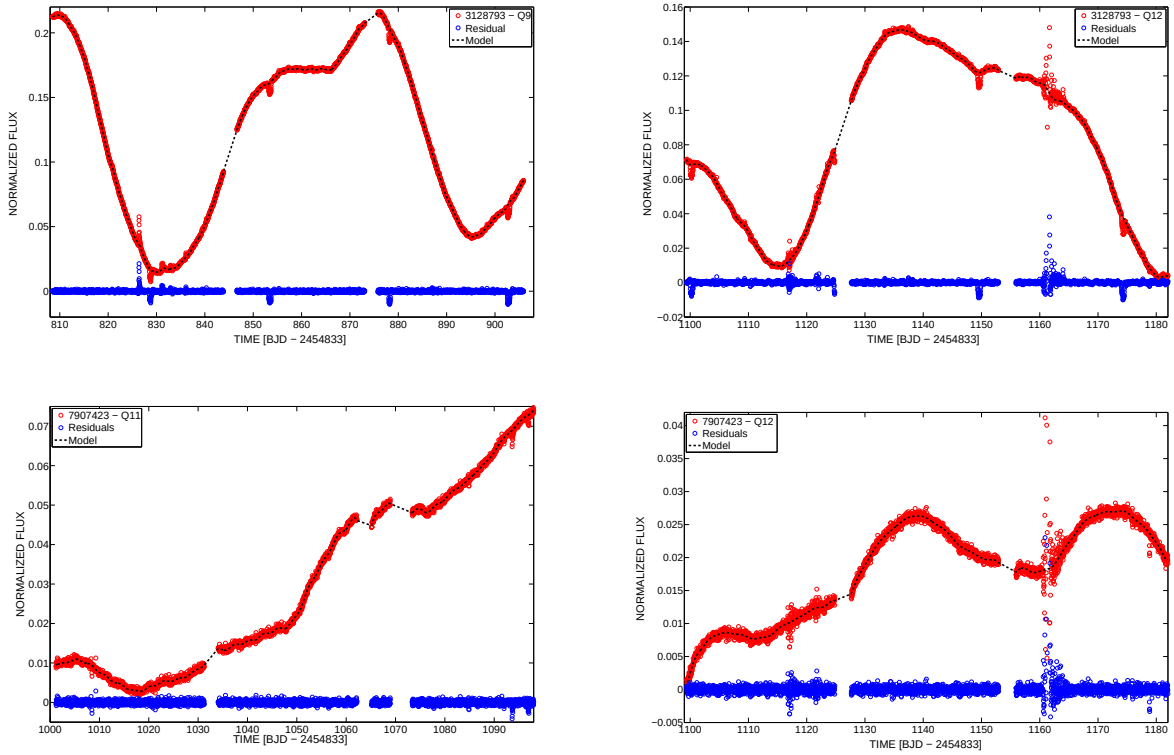

Fig. 8.— Fit to the raw data (red circles) and respective residuals (blue circles) after the GP detrend (dotted line). We plotted a couple of quarters for each analysed planetary candidate (from top to bottom: KIC 3128793 (Q 9,12); KIC 7907423 (Q 11,12)). The raw data are shifted for clarity. Note that both of the datasets on the right-hand side show a large intensity variation (beyond 5σ) immediately after epoch 1160, probably due to problems with the telescope. Nevertheless our GP technique is not affected by these outliers and it catches the main stellar trend, relocating the large variation in the residuals.

REFERENCES

Ballard, S., Fabrycky, D., Fressin, F., et al. 2011, ApJ, 743, 200

Barclay, T., Rowe, J. F., Lissauer, J. J., et al. 2013, Nature, 494, 452

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977

Borucki, W. J., Koch, D. G., Basri, G., et al. 2011, ApJ, 736, 19

Borucki, W. J., Agol, E., Fressin, F., et al. 2013, DOI: 10.1126/science.1234702

Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, MNRAS, 429, 165

Christiansen, J. L., et al. 2012, Kepler Data Characteristics Handbook (KSCI-19040-003)

Ford, E. B., Fabrycky, D. C., Steffen, J. H., et al. 2012, ApJ, 750, 113

García, R. A., Hekker, S., Stello, D., et al. 2011, MNRAS, 414, L6

Gibson, N. P., Aigrain, S., Roberts, S., et al. 2012, MNRAS, 419, 2683

Gilliland, R. L., Jenkins, J. M., Borucki, W. J., et al. 2010, ApJ, 713, L160

Haario, H., Laine, L., Mira, A., Saksman, E., 2006, Statistics and Computing, 16, 339

Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010a ApJ, 713, L87

Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010b, ApJ, 713, L120

Mandel, K., & Agol, E. 2002, ApJ, 580, L171

Murphy, S. J. 2012, MNRAS, 422, 665

Quiñonero, Candela and Rasmussen *A unifying view of sparse approximate Gaussian process regression.* Journal of Machine Learning Research, 2005

Rasmussen, C. and Williams, C. . *Gaussian Processes for Machine Learning*. MIT Press, 2006.

Seikel, M., Clarkson, C., & Smith, M. 2012, J. Cosmology Astropart. Phys., 6, 36

Smith, J. C., Stumpe, M. C., Van Cleve, J. E., et al. 2012, PASP, 124, 1000

Steffen, J. H., Fabrycky, D. C., Agol, E., et al. 2013, MNRAS, 428, 1077

Thatte, A., Deroo, P., & Swain, M. R. 2010, A&A, 523, A35

Van Cleve, J. E., & Caldwell,D. A., Kepler Instrument Handbook (KSCI-19033)

Waldmann, I. P. 2012, ApJ, 747, 12

Waldmann, I. P. 2013, arXiv:1302.6714

Waldmann, I. P., Tinetti, G., Deroo, P., et al. 2013, ApJ, 766, 7